

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 730 230 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
04.09.1996 Bulletin 1996/36

(51) Int. Cl.⁶: G06F 11/14, G06F 11/00

(21) Application number: 96103124.2

(22) Date of filing: 01.03.1996

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE

(30) Priority: 03.03.1995 US 397910

(71) Applicant: HAL COMPUTER SYSTEMS, INC.
Campbell, CA 95008 (US)

(72) Inventors:
• Chang, Chih-Wei David
Saratoga, CA 95070 (US)

• Boney, Joel Frederick
Cupertino, CA 95014 (US)
• Kohli, Jaspal
Sunnyvale, CA 94087 (US)

(74) Representative: Liesegang, Roland, Dr.-Ing.
FORRESTER & BOEHMERT
Franz-Joseph-Strasse 38
80801 München (DE)

(54) Method and apparatus for prioritizing and handling errors in a computer system

(57) A computer system (10) includes a central processing unit (12) and a memory management unit (18) having a plurality of functional units, such as a memory interface unit, a remote interface unit (60), a cache interface unit (70), and a translation unit (50). Each functional unit has a low priority error queue for storing error information for errors having a low priority. Some functional units also have a high priority error queue for storing error information for errors having a high priority error. Based on the status of the error queues, the memory management unit prioritizes and

handles errors caused by hardware failures. For low priority errors, an interrupt request signal is sent to the central processing unit (122). For high priority errors, a RED ALERT signal is sent to the processing unit (112) to cause the processing unit to give immediate attention to the error. For high priority error queue overflows, a failure signal is generated (102) which causes the system to be halted and the contents of the system to be scanned out (104). Thus, errors are prioritized and handled accordingly.

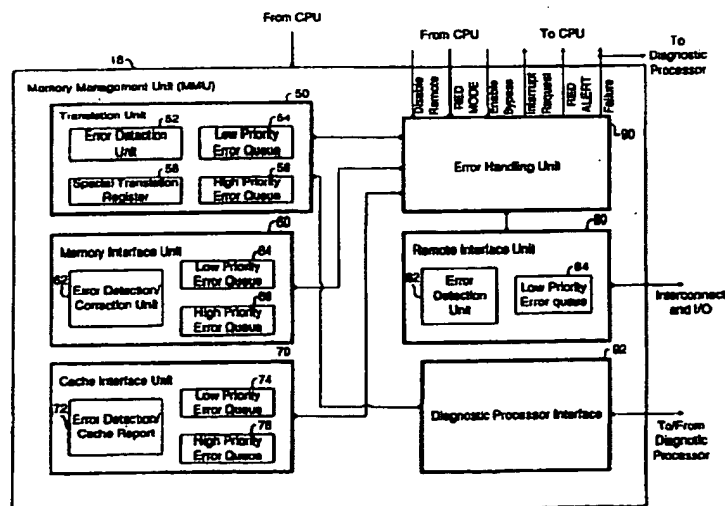


Fig. 2

Description

Related Applications

The subject matter of this application is related to the subject matter of the following applications:

European patent application 96101842.1;

European patent application 96101839.7;

European patent application 96101840.5;

European patent application 96101841.3;

the European patent application entitled "METHOD AND APPARATUS FOR ACCELERATING CONTROL TRANSFER RETURNS";

the European patent application entitled "METHOD AND APPARATUS FOR SELECTING INSTRUCTIONS FROM ONES READY TO EXECUTE";

the European patent application entitled "METHODS FOR UPDATING FETCH PROGRAM COUNTER";

the European patent application entitled "METHOD AND APPARATUS FOR RAPID EXECUTION OF CONTROL TRANSFER INSTRUCTIONS";

the European patent application entitled "ECC PROTECTED MEMORY ORGANIZATION WITH PIPELINED READ-MODIFY-WRITE ACCESSES";

the European patent application entitled "RECLAMATION OF PROCESSOR RESOURCES IN A DATA PROCESSOR";

the European patent application entitled "HARDWARE SUPPORT FOR FAST SOFTWARE EMULATION OF UNIMPLEMENTED INSTRUCTIONS"; and

the European patent application entitled "METHOD AND APPARATUS FOR GENERATING A ZERO BIT STATUS FLAG IN A MICROPROCESSOR",

the latter eight of which are filed simultaneously with this application.

Field of the Invention

This invention relates generally to computer systems and more particularly to a method and apparatus for prioritizing and handling hardware errors in a computer system.

Background of the Invention

In recent years, computer systems have progressively become larger and more complex. The larger a computer system is, the more components it contains, and the more components there are, the greater the chances of hardware failure. As a result, for very large and complex computer systems, hardware failures are practically inevitable. Since hardware failure is almost a given, the important issue in large-scale computer systems becomes the manner in which hardware failures or errors are handled.

Hardware failures fall into several different categories. A first category is that of correctable failure. For this type of failure, operation of the computer system need not be immediately interrupted since the error can be

corrected. A second category is that of non-correctable error. With this type of failure, system operation is immediately interrupted in order to prevent the system from using corrupted data or executing a corrupted instruction. This type of hardware failure typically causes the system to re-execute an instruction or to repeat a particular process. A third type of hardware failure is one in which there is no possibility of recovery. With this type of failure, the system needs to be shut down and restarted. As can be seen from this discussion, the different categories of hardware failures require different handling.

In order to maximize system efficiency, hardware failures should be prioritized and handled accordingly. Currently, however, there is no system believed to be available which carries out this function satisfactorily and efficiently.

Summary of the Invention

In accordance with the present invention, there is provided a computer system wherein hardware failures are efficiently prioritized and handled. In the preferred embodiment, the computer system comprises a central processing unit (CPU), at least one cache, and a memory management unit (MMU) wherein a plurality of low priority and high priority error queues are maintained. Each queue is associated with a selected unit of the MMU. Whenever a low priority error (e.g. a correctable error) is detected in one of the MMU units, an entry is loaded into the low priority queue associated with that MMU unit. Once loaded with an entry, the low priority queue sends out a control signal indicating that a low priority error has occurred. In response, the MMU sends an interrupt request signal to the CPU. Depending on the level of the interrupt request (which may be set by a user) and the status of a mask register within the CPU (which may also be set by a user), the interrupt may either be serviced by the CPU or it may be ignored for the time being. Regardless of which action is taken by the CPU, system operation continues because the error is correctable. Primarily, entries in the low priority error queues are used for purposes of logging the hardware failure for subsequent analysis.

On the other hand, if a high priority error (e.g. a non-correctable error) is encountered by one of the MMU units, then an entry is loaded into the high priority error queue associated with that MMU unit. Once that is done, the high priority queue sends out a control signal indicating that a non-correctable error has been detected. In response, the MMU sends a RED ALERT control signal to the CPU to cause the CPU to give immediate attention to the error. Thus, a non-correctable error is given much higher priority than a correctable error. In general, non-correctable errors may cause termination of the currently executing instruction or program but it usually does not necessitate halting the whole system.

Finally, it may be possible that one or more of the high priority error queues may overflow, thereby indicat-

ing that more non-correctable errors have been detected than the system can handle. If this happens, then one or more of the high priority queues will issue an overflow signal. In response to this overflow signal, the MMU will issue a control signal to stop the system clock. This serves to freeze the system at the current state. Thereafter, the contents of the system are scanned out to ascertain the internal states of the system. This process is preferably carried out only when it becomes clear that recovery from non-correctable errors or failures is not possible, i.e. when one or more of the high priority queues overflows.

As shown by the above discussion, the present invention prioritizes hardware failures based on the type of hardware error. In addition, each type of failure is handled in an efficient manner suitable for the type of error. Overall, the present invention provides an efficient and effective means for prioritizing and handling hardware failures.

Brief Description of the Drawings

Fig. 1 is a block diagram representation of a computer system 10 wherein the present invention is implemented.

Fig. 2 is a more detailed block diagram of the memory management unit 18 of the present invention.

Fig. 3 is a flow diagram for the error handling unit 90 of Fig. 2.

Detailed Description of the Preferred Embodiments

With reference to Fig. 1, there is shown a computer system 10 wherein the present invention is implemented, the system 10 preferably comprising a central processing unit (CPU) 12, an instruction cache 14 for storing recently executed instructions, a data cache 16 for storing recently accessed data, a memory 20, a memory management unit (MMU) 18 for coordinating access to the memory 20, and a clock unit 22. System 10 preferably also comprises a diagnostic processor 24, a random access memory (RAM) 25, a read-only-memory (ROM) 26, and a scan engine 28. As will be explained later, components 24-28 are used for error handling purposes. In the preferred embodiment, the CPU 12 preferably takes the form of a superscalar processor capable of executing a plurality of instructions simultaneously. It should be noted, though, that CPU 12 is not required to be superscalar. Other types of CPU may also be used.

In system 10, normal operational flow is as follows. The CPU 12 initiates operation by generating a virtual address. This virtual address is compared with the address tags stored within the instruction and data caches 14, 16. If a "hit" is found, then the data or instruction is fetched from the caches 14, 16. On the other hand, if a "miss" is encountered, then the virtual address is passed on to the MMU 18 for processing. Upon receiving the virtual address, the MMU 18

responds by translating the virtual address into an address which can be used to access the memory 20, and then fetching the instruction or data from the memory 20. Thereafter, the requested data or instruction is passed on to the CPU 12 for processing.

In general, MMU 18 of system 10 performs five major functions. First, MMU 18 translates virtual addresses from the CPU 12 into addresses which can be used to access the memory 20. Second, MMU 18 provides an interface to the memory 20 for accessing and retrieving information therefrom. Third, MMU 18 provides an interface to the caches 14, 16 so that when information is retrieved from the memory 20, the information is stored into one of the caches. Fourth, MMU 18 provides an interface to the interconnected system (i.e. bus system) and input/output (I/O) devices. This interface is used, for example, to control direct memory access (DMA) between an external device and the memory 20. In addition to the previous functions, MMU 18 preferably further performs the error prioritization and handling function of the present invention. This function will be described in greater detail in a subsequent section.

The MMU 18 is shown in greater detail in Fig. 2. As shown, MMU 18 preferably comprises a translation unit 50, a memory interface unit 60, a cache interface unit 70, a remote interface unit 80, an error handling unit 90, and a diagnostic processor interface 92. With regard to translation unit 50, it is this unit 50 which translates or maps the virtual addresses received from the CPU 12 into addresses which can be used to access the memory 20. In the preferred embodiment, unit 50 comprises an error detection unit 52 for detecting possible translation errors, a low priority error queue 54 for storing low priority error information, a high priority error queue 56 for storing high priority error information, and a special translation register 58 for storing an address translation used in the error handling process. Preferably, each of the error queues 54, 56 contains a plurality of entries so that more than one set of error information can be stored in each.

In normal operation, translation unit 50 receives and translates virtual addresses from the CPU using translation tables (not shown) within the translation unit 50. In the course of carrying out this translation function, the error detection unit 52 of unit 50 checks the address translations for possible errors caused by hardware failures. If a low priority error (i.e. an error which does not require immediate attention from the CPU 12, such as a single bit hardware correctable error) is detected, then the error is logged into an entry of the low priority error queue 54. Preferably, the information stored in queue 54 includes specific error information such as the type of error, where the error occurred, and information relating to the nature of the error. If one or more entries are logged into the low priority error queue 54, then queue 54 will send a low priority error signal to the error handling unit 90.

On the other hand, if a high priority error (i.e. an error which prevents the current access from being completed, such as a multiple bit non-correctable error) is detected, then an entry is entered into the high priority error queue 56. This entry preferably includes specific error information such as error type, location of error, and information relating to the nature of the error. If one or more entries are loaded into the high priority error queue 56, then queue 56 will send a high priority error signal to the error handling unit 90. As an additional function, the high priority error queue 56 preferably generates and sends an overflow signal to the error handling unit 90 if an attempt is made to write an entry to the queue 56 when the queue 56 is full. This overflow signal indicates to the error handling unit 90 that more errors have been encountered than the queue 56 can handle. As will be explained in a subsequent section, the low priority error signal, the high priority error signal, and the overflow signal are processed by the error handling unit 90 to determine the proper course of action.

With regard to the memory interface unit 60, it is this unit 60 which takes the translated addresses from the translation unit 50 and uses the translated addresses to access the memory 20 to retrieve information therefrom. Memory interface unit 60 preferably comprises an error detection/correction unit 62, a low priority error queue 64, and a high priority error queue 66. Queues 64 and 66 are substantially identical to queues 54 and 56 of the translation unit 50. In performing the interfacing function, the error correction unit 62 of unit 60 checks information from the memory 20 for possible errors caused by hardware failures. If a low priority error such as a single bit hardware correctable error is detected, then detection/correction unit 62 preferably corrects the error and thereafter logs the error into an entry of the low priority queue 64. One or more entries in the low priority error queue 64 will cause the queue 64 to send a low priority error signal to the error handling unit 90. If instead a high priority error such as a multiple bit non-correctable error is detected, then unit 62 preferably writes an entry into the high priority error queue 66. One or more entries in the high priority error queue 66 causes the queue 64 to send a high priority signal to the error handling unit 90. In addition, if the error detection/correction unit 62 attempts to write an entry into queue 64 when the queue is already full, then queue 64 generates and sends an overflow signal to the error handling unit 90.

The cache interface unit 70 of MMU 18 is the unit which handles the exchange of information between the caches 14, 16 and the MMU 18. More specifically, the cache interface unit 70 handles the loading of information retrieved from memory 20 into the caches 14, 16, and the storing of information from the caches 14, 16 into memory 20. Interface unit 70 preferably comprises an error detection/cache report unit 72, a low priority error queue 74, and a high priority error queue 76. Queues 74 and 76 are preferably substantially identical to queues 54 and 56 of the translation unit 50. In the

preferred embodiment, the caches 14, 16 preferably comprise mechanisms for detecting and correcting (if possible) the errors caused by hardware failures within the caches 14, 16; thus, unit 74 preferably does not perform this function. However, errors are preferably reported by the caches 14, 16 to unit 72 of the interface unit 70. In response, unit 72 preferably makes a determination with regard to the error reported. If the error is a low priority error, such as a single bit correctable error, then error information is written into an entry of the low priority error queue 74. Writing one or more entries into queue 74 causes the queue 74 to send a low priority signal to the error handling unit 90. On the other hand, if the error is a high priority error, such as a multiple bit non-correctable error, then unit 72 writes error information into the high priority error queue 76. Writing one or more entries into queue 76 causes the queue to send a high priority error signal to the error handling unit 90. In addition, queue 76 preferably generates and send an overflow signal to the error handling unit 90 if unit 72 attempts to write an entry into queue 76 when the queue is already full.

MMU 18 preferably further comprises a remote interface unit 80 for interacting with an interconnect system and the I/O devices coupled thereto. It is unit 80 which, for example, controls DMA access to the memory 20 by an I/O device. Preferably, unit 80 comprises an error detection unit 82 for detecting low priority errors coming from the interconnect system, and a low priority error queue 84. If a low priority error is detected, then unit 82 writes error information into an entry of queue 84. Writing one or more entries into queue 84 causes the queue to send a low priority error signal to the error handling unit 90.

The error handling unit 90 and the diagnostic processor interface 92 are the two units on the MMU 18 which are responsible for coordinating the prioritization and handling of errors. Preferably, error handling unit 90 receives all of the low priority error signals, high priority error signals, and overflow signals from all of the units 50, 60, 70, and 80. Armed with this information, unit 90 determines which course of action to take with regard to error prioritization and handling. Fig. 3 shows an operational flow diagram for error handling unit 90. Preferably, unit 90 begins operation by checking 100 for an overflow signal from one of the high priority error queues 56, 66, 76. If an overflow signal is detected, then it means that at least one of the units 50, 60, 70 has encountered more high priority errors or failures than it can handle. In such a case, the system 10 should be halted. To accomplish this, error handling unit 90 first sends 102 a failure signal to the clock unit 22 (Fig. 1). This serves to freeze the current state of the system 10. In addition, unit 90 sends 104 the failure signal to the diagnostic processor 24 (via scan engine 28) to inform the processor 24 that system failure has been experienced. In response, diagnostic processor 24 accesses and executes a scan control program 32 stored within the ROM 26. Under control of program 32, processor 24 interacts with the scan

engine 28 to scan out the contents of the system components 12, 14, 16, 18. By so doing, the state of the system 10 is saved so that it may be later analyzed to determine the cause of the system failure.

Returning to step 100, if none of the overflow signals from error queues 56, 66, and 76 are asserted, then error handling unit 90 goes on to check 110 the status of the high priority error signals from the high priority error queues 56, 66, 76. If any one of these error signals is asserted, then it means that an error has occurred which requires the immediate attention of the CPU 12. In such a case, error handling unit 90 preferably generates and sends 112 a RED ALERT signal to the CPU 12. In response to this signal, the CPU 12 enters RED MODE, wherein a number of operations are performed. In RED MODE, CPU 12 first puts itself into sequential operation (i.e. processing only one instruction at a time) instead of superscalar operation. Second, CPU 12 invalidates and disables its on-chip cache, and also disables the instruction and data caches 14, 16. In addition, CPU 12 generates and sends several control signals to the MMU 18. These control signals include a RED MODE confirmation signal, a bypass signal, and a disable remote signal.

Upon receiving 114 these control signals from the CPU 12, the error handling unit 90 proceeds to step 116 to disable the remote interface unit 80 by sending a disable signal to the unit 80. This serves to block further I/O bus access by external I/O devices. Also, in step 116, error handling unit 90 enables the bypass feature of the translation unit 50 by sending an enable signal to the unit 50. Once activated, this bypass feature causes the translation unit 50 to deviate from its regular operation. Instead of using regular translation tables to perform its address translations, the translation unit in bypass mode uses the special translation register 58 to perform address translation. Preferably, register 58 contains a single address translation entry. After step 118 is performed, the system 10 is ready for RED MODE operation.

RED MODE operation preferably begins with the CPU 12 issuing a request for an instruction, the request preferably including a specific virtual address and a load command. This request is sent to the MMU 18, and more specifically, the virtual address is sent to the translation unit 50 and the command is sent to the diagnostic processor interface 92. In response, the translation unit 50 uses the special translation register 58 to provide a translated address for the virtual address. This translated address is sent to the diagnostic processor interface 92. In response, the diagnostic processor interface 92 sends the translated address and the load command to the diagnostic processor 24 for processing.

Upon receipt of the load command and the translated address, the diagnostic processor 24 processes the load command to retrieve information from the ROM 26 from a location indicated by the translated address. Preferably, the ROM 26 contains therein a section 34 wherein RED MODE code is stored, and preferably the

translated address points to a location within section 34. By processing the load command, the diagnostic processor 24 is in effect retrieving a RED MODE instruction from the ROM 26 for the CPU 12 to execute. Once the instruction is retrieved, it is passed on to the diagnostic processor interface 92, which in turn, passes the instruction on to the CPU 12 for execution. Armed with this RED MODE instruction, the CPU 12 can begin executing RED MODE code to properly process the high priority errors. Preferably, the CPU 12 continues this process of fetching RED MODE code by way of the diagnostic processor 24 as long as RED MODE is invoked.

Under control of the RED MODE code, CPU 12 preferably processes the high priority error or errors by reading the high priority error queues 56, 66, 76. For each high priority error found in the queues 56, 66, 76, CPU 12 preferably carries out a proper procedure to rectify or to circumvent the error. The specific procedure carried out by the CPU 12 will depend on the nature of the error and the specific configuration of the system, and thus, is application-specific. Preferably, once CPU 12 is in RED MODE, it processes all of the high priority errors in the high priority error queues 56, 66, 76 before exiting RED MODE. Once an error is rectified, the corresponding entry in the high priority error queue is cleared. High priority errors are thus handled.

Returning to step 110, if none of the high priority error signals from queues 56, 66, 76 are asserted, then error handling unit 90 proceeds to step 120 to determine whether any of the low priority error signals from the low priority error queues 54, 64, 74, 84 are asserted. If one or more of these low priority error signals is asserted, then error handling unit 90 will generate 122 and send an interrupt request signal to the CPU 12 to inform the CPU 12 that a low priority error has occurred. The level of this interrupt request can be set by a user. Also, within the CPU 12, there is an interrupt mask register 30. The contents of this register 30, which can also be set by the user, are used to mask out certain interrupt signals. Depending on the interrupt level of the interrupt request, and the contents of the mask register 30, the CPU 12 may or may not process the interrupt immediately. If the CPU 12 does not service the interrupt, then error handling unit 90 preferably maintains the interrupt request signal in the active state. Operation of MMU 18 remains the same. Low priority errors continue to be stored into the low priority error queues 54, 64, 74. If these queues become full, then the new entries will simply overwrite the old entries. Since the low priority error entries are used primarily for logging purposes, overwriting some of the error entries will not adversely affect system operation.

If, on the other hand, the CPU 12 decides to service the interrupt, then the error entries in all of the low priority error queues 54, 64, 74 are read, processed, and then cleared by the CPU 12. Once that is done, the interrupt signal is deasserted and the system 10 returns to normal operation.

The present invention has been described with reference to a specific embodiment. It should be noted, however, that the invention should not be construed to be so limited. Various modifications may be made by one of ordinary skill in the art with the benefit of this disclosure without departing from the spirit of the invention. Therefore, the present invention should not be limited by the examples used to illustrate it but only by the scope of the appended claims.

5

10

Claims

1. A method for handling memory errors in a computer having a memory, the computer operating responsive to a clock, the method comprising the steps of:
 - detecting the occurrence of a memory error;
 - identifying the type of memory error as either a first type or a second type;
 - storing in a first error queue an address of the memory error if the error is a first type of error;
 - storing in a second error queue an address of the memory error if the error is a second type;
 - detecting an overflow if more than a predetermined number of addresses are stored in the second error queue;
 - disabling the clock responsive to the detected overflow.

15

20

25

30

35

40

45

50

55




Fig. 1

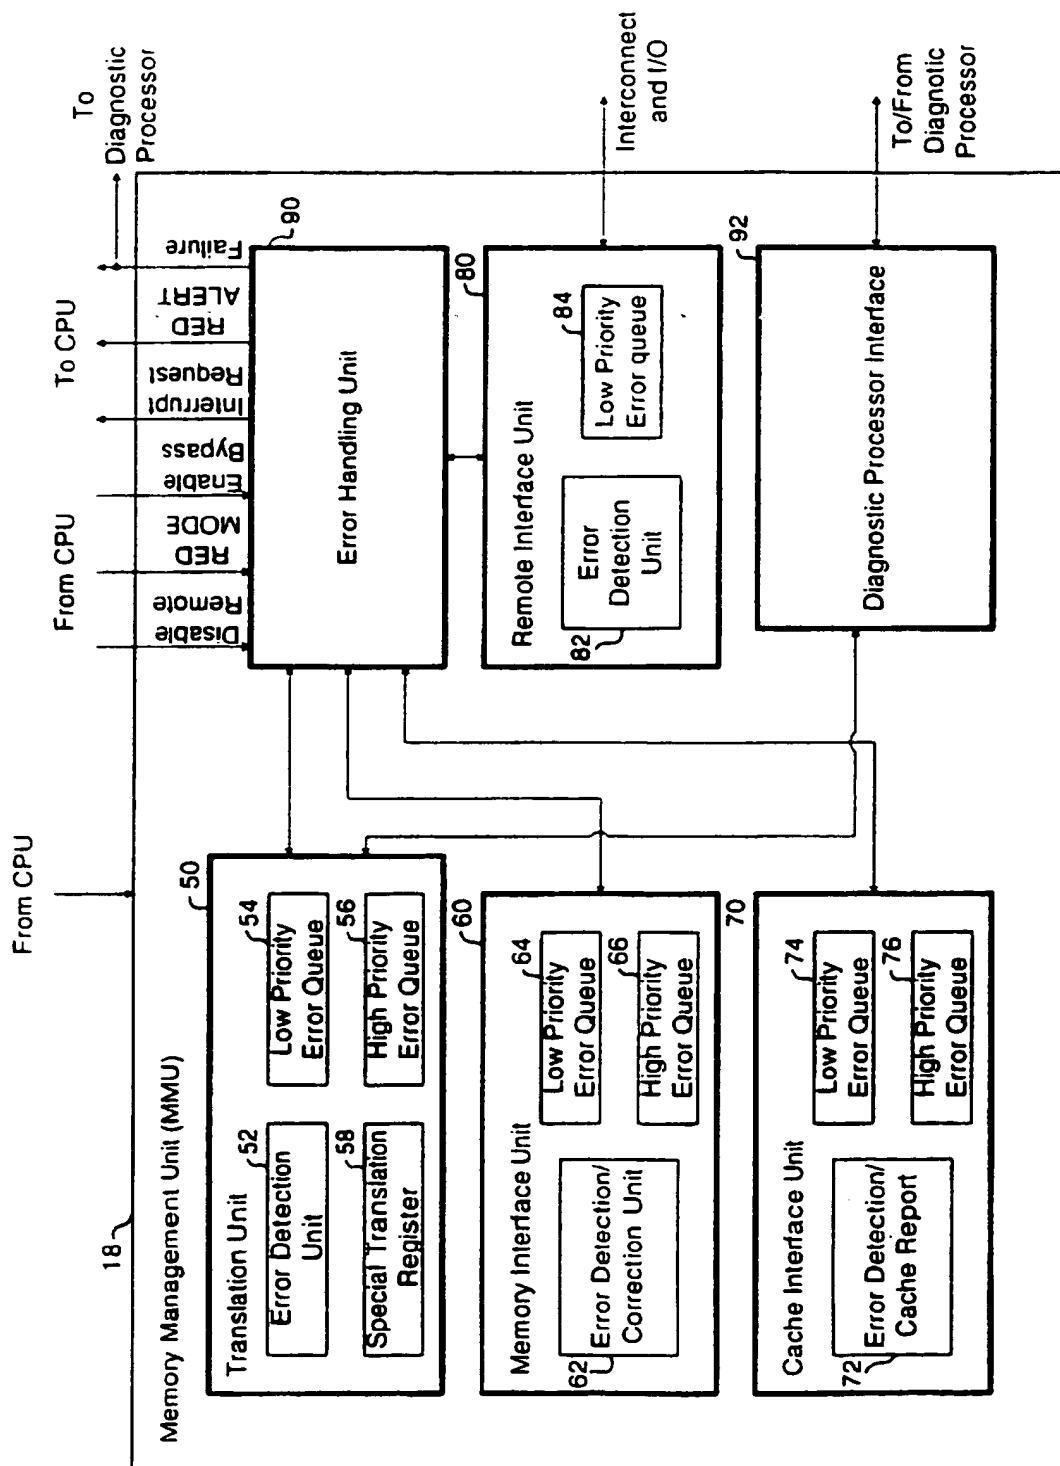


Fig. 2




Fig. 3




Fig. 1

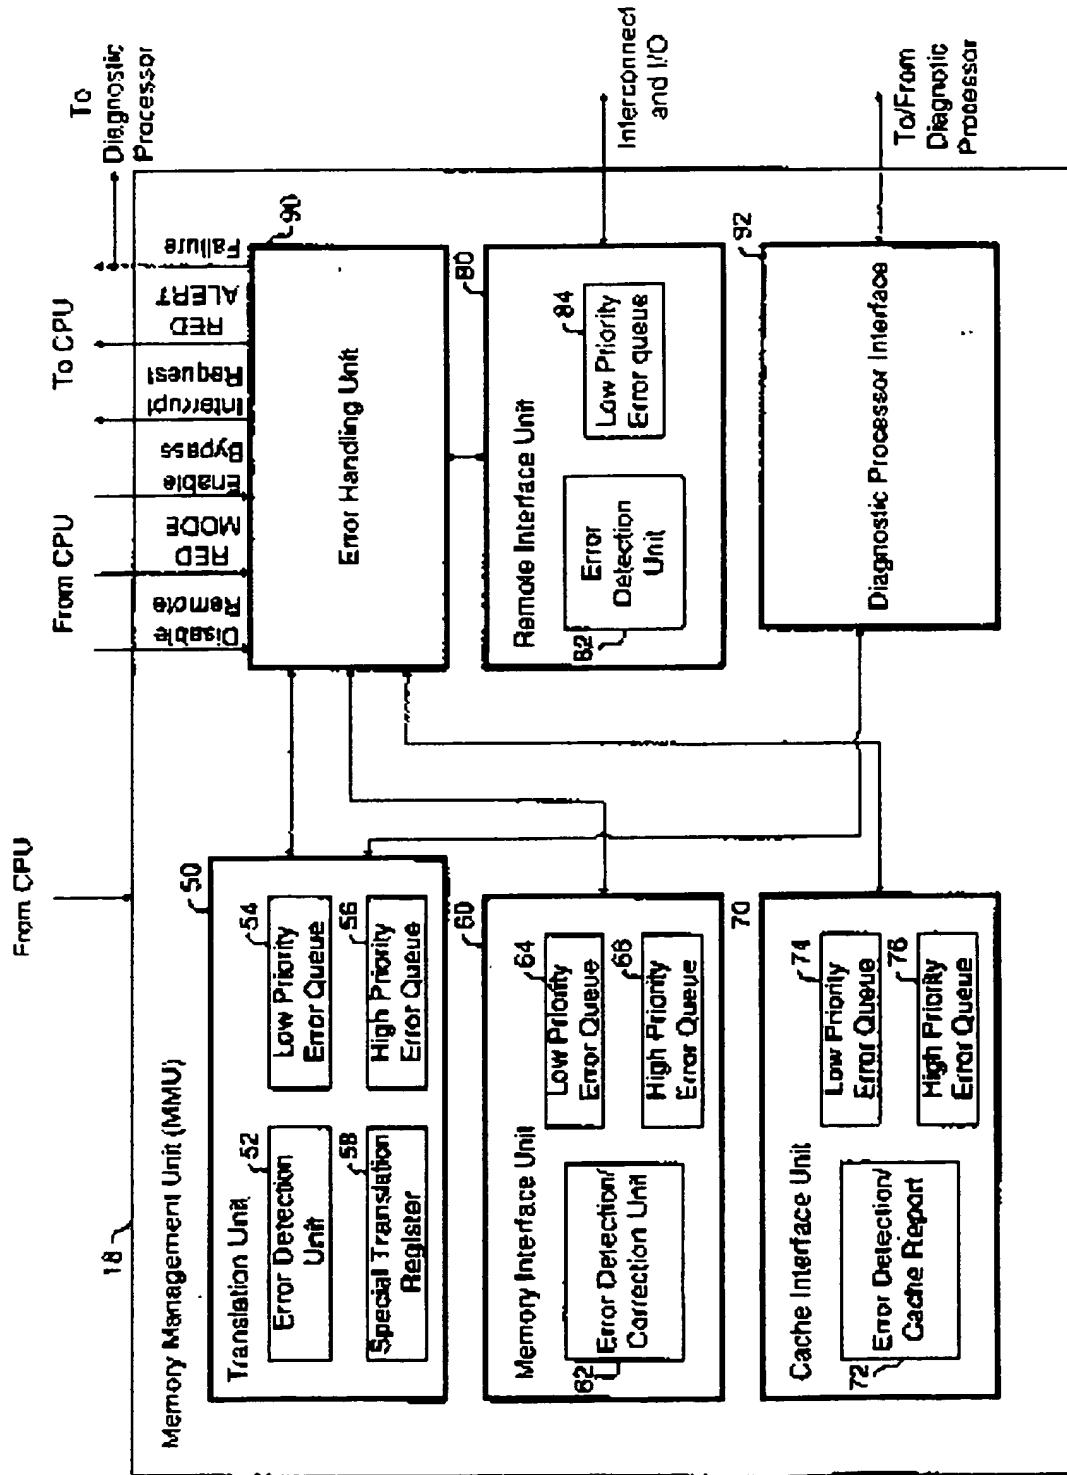


Fig. 2





Fig. 3



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 730 230 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
22.12.1999 Bulletin 1999/51

(51) Int. Cl.⁶: G06F 11/14, G06F 11/00

(43) Date of publication A2:
04.09.1996 Bulletin 1996/36

(21) Application number: 96103124.2

(22) Date of filing: 01.03.1996

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE

(30) Priority: 03.03.1995 US 397910

(71) Applicant:
HAL COMPUTER SYSTEMS, INC.
Campbell, CA 95008 (US)

(72) Inventors:
• Chang, Chih-Wei David
Saratoga, CA 95070 (US)

• Boney, Joel Frederick
Cupertino, CA 95014 (US)
• Kohli, Jaspal
Sunnyvale, CA 94087 (US)

(74) Representative:
Liesegang, Roland, Dr.-Ing.
FORRESTER & BOEHMERT
Franz-Joseph-Strasse 38
80801 München (DE)

(54) Method and apparatus for prioritizing and handling errors in a computer system

(57) A computer system (10) includes a central processing unit (12) and a memory management unit (18) having a plurality of functional units, such as a memory interface unit (60), a remote interface unit (80), a cache interface unit (70), and a translation unit (50). Each functional unit has a low priority error queue for storing error information for errors having a low priority. Some functional units also have a high priority error queue for storing error information for errors having a high priority error. Based on the status of the error queues, the memory management unit prioritizes and handles errors caused by hardware failures. For low priority errors, an interrupt request signal is sent to the central processing unit (122). For high priority errors, a RED ALERT signal is sent to the processing unit (112) to cause the processing unit to give immediate attention to the error. For high priority error queue overflows, a failure signal is generated (102) which causes the system to be halted and the contents of the system to be scanned out (104). Thus, errors are prioritized and handled accordingly.

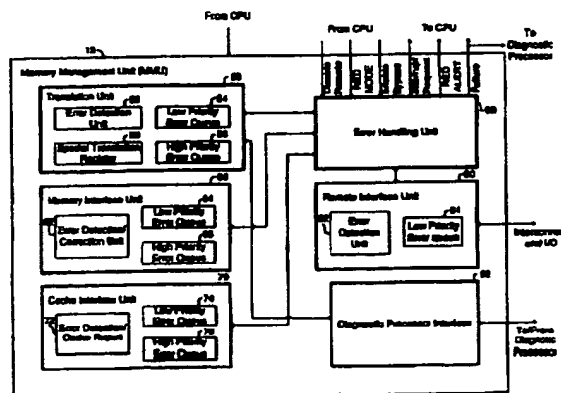


Fig. 2

EP 0 730 230 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 10 3124

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	US 4 866 712 A (CHAO MING-TE) 12 September 1989 (1989-09-12) * column 2, line 7 - line 52 * * column 4, line 52 - column 5, line 43; figure 3 *	1	G06F11/14 G06F11/00
A	EP 0 335 507 A (DIGITAL EQUIPMENT CORP) 4 October 1989 (1989-10-04) * column 4, line 16 - line 26 *	1	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 18 October 1999	Examiner Fernandez Balseiro, J
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 96 10 3124

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

18-10-1999

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 4866712 A	12-09-1989	AU 4031089 A	06-09-1989
		CA 1310422 A	17-11-1992
		EP 0402403 A	19-12-1990
		WO 8907795 A	24-08-1989
EP 0335507 A	04-10-1989	AU 3088889 A	05-10-1989
		CA 1318030 A	18-05-1993
		DE 68924923 D	11-01-1996
		DE 68924923 T	25-07-1996
		JP 2016640 A	19-01-1990
		JP 2019495 C	19-02-1996
		JP 7058474 B	21-06-1995
		US 5090014 A	18-02-1992
		US 5469463 A	21-11-1995